



ANALYSIS OF SMALL SCALE MBR ANISOTROPY IN THE PRESENCE OF FOREGROUND CONTAMINATION ¹

SCOTT DODELSON and ALBERT STEBBINS

NASA/Fermilab Astrophysics Center
Fermi National Accelerator Laboratory
P.O. Box 500, Batavia, IL 60510
USA

dodelson@virgo.fnal.gov

stebbins@fnalv.fnal.gov

Abstract

Many of the current round of experiments searching for anisotropies in the Microwave Background Radiation (MBR) are confronting the problem of how to disentangle the cosmic signal from contamination due to galactic and intergalactic foreground sources. Here we show how commonly used likelihood function techniques can be generalized to account for foreground. Specifically we set some restrictions on the spectrum of foreground contamination but allow the amplitude to vary arbitrarily. The likelihood function thus generalized gives reasonable limits on the MBR anisotropy which, in some cases, are not much less restrictive than what one would get from more detailed modeling of the foreground. Furthermore, the likelihood function is exactly the same as one would obtain by simply projecting out foreground contamination and just looking at the reduced data set. We apply this generalized analysis to the recent medium angle data sets of ACME-HEMT (Gaier *et al.* 1992, Schuster *et al.* 1993) and MAX (Meinhold *et al.* 1993, Gunderson *et al.* 1993). The resulting analysis constrains the one free parameter in the standard cold dark matter theory to be $Q_{\text{rms-ps}} = 18_{-5}^{+8} \mu K$. This best fit value, although in striking agreement with the normalization from COBE, is not a very good fit, with an overall $\chi^2/\text{degree of freedom} = 208/168$. We also argue against three commonly used methods of dealing with foreground: (i) ignoring it completely; (ii) subtracting off a best fit foreground and treating the residuals as if uncontaminated; and (iii) culling data which appears to be contaminated by foreground.

¹ Submitted to *The Astrophysical Journal*



1. Introduction

Since the detection of Microwave Background Radiation (MBR) anisotropies by the Differential Microwave Radiometer (DMR) on the COsmic Background Explorer (COBE) satellite (Smoot *et al.* 1992), there has been a spate of announcements of detections of anisotropies in microwave brightness on smaller angular scales (Gaier *et al.* 1993, Meinhold *et al.* 1993, Schuster *et al.* 1993, Gunderson *et al.* 1993, and more). Many of these experiments utilize measurements at multiple frequencies in order to be able distinguish primordial anisotropies in the MBR from other types of microwave emission, such as dust, free-free, or synchrotron, which may occur either in our Galaxy or in extra-Galactic objects. While the large scale anisotropies observed by the COBE-DMR appear to have relatively little foreground contamination, the spectrum of many of the small scale anisotropies found are not very well fit by frequency-independent brightness temperature fluctuations and are therefore probably not completely due to primordial anisotropy. As stressed by Brandt *et al.* , (1993) it will take more and better observations to be able to disentangle the primordial anisotropies from the foreground contamination. In the meantime one would like to use the small-scale measurements to set limits on the parameters of models of cosmological inhomogeneities. To obtain such limits one must take into account the uncertain contamination of the measured anisotropies. In this paper we will discuss some methods of accounting for the contamination when constraining parameters and apply them to some of the ACME-HEMT (Gaier *et al.* 1992, Schuster *et al.* 1993) and MAX (Meinhold *et al.* 1993, Gunderson *et al.* 1993) data. We stress that taking into account unknown amounts of contamination involves great uncertainties and different approaches can be expected to yield different limits on parameters. Here the emphasis will be on setting reliable, and perhaps conservative, limits on model parameters.

In §2 we introduce the likelihood function which is the standard tool used to compare theoretical predictions with MBR anisotropy data. We then go on to discuss various properties of the generalization of the likelihood function to include a known statistical distribution of sources. We show that this generalized likelihood function has a well-defined and useful limit when we take the amplitude of the foreground anisotropies to be large. This limit is independent of the spatial correlations of the foreground emission and depends only on the assumed spectra of the various components of foreground contamination. This limit is also equivalent to “projecting out” or “marginalizing” (Anthony Lasenby’s terminology) the foreground emission from multifrequency data. Finally in §2 we show how this limit may be easily taken when the MBR model and detector noise are assumed Gaussian. In §3 we give some examples of how the procedure behaves by applying it to the data of Gaier *et al.* (1993). In §4 we apply it to rest of the data. Finally in §5 we discuss the meaning of the results obtained.

2. Likelihood Functions

There are at least three components to microwave temperature anisotropies measured on the sky. For cosmologists the most interesting component is the set of primordial MBR anisotropies. These are caused by the initial inhomogeneities in the Universe which are thought to be set up by some random process occurring at very early cosmological times. The particular random process we will refer to as a cosmological model. Here we will suppose one is considering only a subset of possible models, parameterized by a few numbers. Let us denote these numbers by the shorthand \mathcal{P} , for parameters. What we would like to do is estimate \mathcal{P} from a set of MBR anisotropy measurements. Let \mathcal{T} represent the *true* values of the primordial MBR anisotropies that have been estimated by experiments. Given \mathcal{T} we can estimate \mathcal{P} using the probability density of \mathcal{T} given \mathcal{P} , i.e. $p_1(\mathcal{T}|\mathcal{P})d\mathcal{T}$. In addition to the primordial MBR anisotropy there is also a component of foreground contamination which adds to the signal coming into an MBR experiment, or symbolically $\mathcal{S} = \mathcal{T} + \mathcal{F}$, where \mathcal{S} and \mathcal{F} refer to external signal and foreground contamination. The third ingredient is the observational error or detector noise, \mathcal{N} , in measuring brightness fluctuations. Thus the data we obtain, \mathcal{D} , can be written as $\mathcal{D} = \mathcal{S} + \mathcal{N} = \mathcal{T} + \mathcal{F} + \mathcal{N}$. By testing and good experimental design we usually know the distribution of \mathcal{N} or mathematically $p_2(\mathcal{D}|\mathcal{S})d\mathcal{D}$ which is the probability density of \mathcal{D} given \mathcal{S} . Unfortunately we do not know much about the properties of \mathcal{F} and this limits what we can say about \mathcal{P} .

For the moment assume \mathcal{F} is zero, as is done in many analyses of MBR experiments. Then one has all of the ingredients to construct the probability density of \mathcal{D} conditional on \mathcal{T} , i.e.

$$p_3(\mathcal{D}|\mathcal{P})d\mathcal{D} = \left[\int p_2(\mathcal{D}|\mathcal{T}) p_1(\mathcal{T}|\mathcal{P}) d\mathcal{T} \right] d\mathcal{D} . \quad (1)$$

One approach to inferring \mathcal{P} from \mathcal{D} is to choose \mathcal{P} such that for that value of \mathcal{P} the observed data, \mathcal{D} , is near [but not too near] the mean of the distribution. Along the same lines one might require that \mathcal{P} be chosen such that \mathcal{D} has a high probability density, when compared with other possible values of \mathcal{D} . These methods may be described as requiring “goodness-of-fit” of the data and involve fixing \mathcal{P} and looking at the probability distribution of \mathcal{D} . Of course what we would really like is the probability distribution of \mathcal{P} . Unfortunately there is nothing in probability theory that would allow us to convert the data into such a probability distribution without further assumptions.

Another way to infer \mathcal{P} from \mathcal{D} involves not the probability of \mathcal{P} , but its likelihood. That is, fix the data at the observed value, \mathcal{D} , and choose \mathcal{P} such that the probability density, $p_3(\mathcal{D}|\mathcal{P})$, be large when compared to other values of \mathcal{P} . Clearly a large probability density near the observed data favors such a value of \mathcal{P} . Another motivation for this procedure stems from Bayes’ Theorem. Let us suppose we had made the further assumption (or had further knowledge) that the

parameters, \mathcal{P} , were determined by some random process, with probability distribution $p_4(\mathcal{P}) d\mathcal{P}$. We wish to evaluate the \mathcal{P} used to generate the data, \mathcal{D} . Bayes' theorem says that given that the measurements yield \mathcal{D} , the new (*posterior*) probability distribution for \mathcal{P} is just the original (*prior*) distribution times the likelihood function, L , or mathematically

$$p_5(\mathcal{P}|\mathcal{D}) d\mathcal{P} = L(\mathcal{P}; \mathcal{D}) p_4(\mathcal{P}) d\mathcal{P} \quad L(\mathcal{P}; \mathcal{D}) \equiv \frac{p_3(\mathcal{D}|\mathcal{P})}{\int p_3(\mathcal{D}|\mathcal{P}) p_4(\mathcal{P}) d\mathcal{P}}. \quad (2)$$

Note that the ratio of L for different values of \mathcal{P} , which is really just the ratio of $p_3(\mathcal{D}|\mathcal{P})$, does not depend on the prior distribution. It is this likelihood ratio which gives the relative increase or decrease of the probability of different values of \mathcal{P} and this is why one might consider values of \mathcal{P} with relatively high likelihood as being preferred. Even if one doesn't have any knowledge of the prior distribution of \mathcal{P} one can still use the likelihood ratio as a statistic to choose \mathcal{P} , although one then cannot determine the probability of the favored choice being correct.

In recent years the likelihood function has been commonly used as a statistic in analyzing anisotropy experiments (see Readhead *et al.* (1989) or Myers (1990) for clear discussions), although usually without regard to any possible foreground contamination. If contamination is present then one must make some assumptions about the contamination in order to construct the likelihood function. If one knows or guesses that the foreground is drawn from some probability distribution, $p_6(\mathcal{F}) d\mathcal{F}$, then

$$L(\mathcal{P}; \mathcal{D}) \propto p_3(\mathcal{D}|\mathcal{P}) = \int \int p_2(\mathcal{D}|T + \mathcal{F}) p_1(T|\mathcal{P}) p_6(\mathcal{F}) dT d\mathcal{F} \quad (3)$$

where we have dropped the \mathcal{P} -independent denominator of L . One could use one's general knowledge of the Galactic dust and gas and of extra-Galactic sources to construct a reasonable p_6 . If available one should make use of measurements of the emission at other frequencies to constrain the foreground emission in the same direction one is measuring the microwave brightness anisotropy. Some radio and infrared data are available for the measurements considered below and we will give some discussion of these in §5. For the moment we will consider the case where we do not have any compelling reason to chose one $p_6(\mathcal{F})$ over another. We now discuss how one can still set believable limits on model parameters even in this state of ignorance.

Marginalization

If we really do not have much idea of what the foreground is doing then to set reliable (= conservative) limits on MBR anisotropy we should take a liberal view of what the foreground emission may do to contaminate the measurements. One might think that such an approach will

lead to no limits at all, but this is not the case. If one sets some restrictions on the spectrum of foreground contamination but allows the amplitude to vary arbitrarily then the likelihood function gives reasonable limits on the MBR anisotropy which, in some cases, are not much less restrictive than what one would get from more detailed modeling of the foreground. Furthermore, in this limit the likelihood function is exactly the same as one would obtain by simply projecting out foreground contamination and just looking at the reduced data set.

Before proceeding to justify these claims we should be a bit more explicit about the meaning of \mathcal{D} , \mathcal{T} , and \mathcal{F} . The brightness pattern on the sky, $I_\nu(\hat{n})$, is a continuous function of both direction, \hat{n} , and frequency, ν . It is the ν dependence which we can use to unambiguously separate MBR anisotropy from foreground contamination. We therefore only consider multifrequency experiments. The experiments only give estimates of the temperature convolved with some window function in direction and frequency, i.e.

$$S = \left\{ \int W_i(\hat{n}) V_{(i,a)}(\nu) I_\nu(\hat{n}) d^2 \hat{n} d\nu, a = 1, \dots, N_{\text{ch}}, i = 1, \dots, N_p \right\}. \quad (4)$$

where N_p gives the number of spatial *patches* and N_{ch} gives the number of spectral *channels* which for simplicity we have assumed is the same for all patches. For simplicity we have also assumed that the window functions factorize into spatial ($W_i(\hat{n})$) and spectral ($V_{(i,a)}(\nu)$) parts and that the spatial window is the same for each channel at that patch. In general we only require that the window function averaged over frequency and weighted by the spectra of any component, i.e. MBR or foreground, be the same for all components and all channels. With this assumption the relation between the signal in different channels of the same patch is telling us only about the spectrum and the detector noise, and not about the spatial pattern of brightness. The moments of both the MBR and foreground brightness which contribute to the measured signal are just the same convolution as given in eq. (4). The MBR brightness in a given direction is given by only one parameter, the temperature $T_{\text{MBR}}(\hat{n})$. We will assume that there are only a finite number of components of contamination, say N_f , each of which have a *known* spectrum. Furthermore we assume that $N_f < N_{\text{ch}}$. If not then any possible observed signal could always be produced by some combination of foreground emission and no MBR anisotropy. Finally \mathcal{D} is just S of eq. (4) with some added detector noise which is different in each patch and channel.

Let us count the number of degrees of freedom (dof) of the various terms which contribute to the $N_p \times N_{\text{ch}}$ observations, \mathcal{D} : \mathcal{T} has N_p dof, \mathcal{F} has $N_p \times N_f$ dof, and \mathcal{N} has $N_p \times N_{\text{ch}}$ dof. This means that there are $(N_{\text{ch}} - N_f) \times N_p$ dof of \mathcal{D} which are linearly independent of \mathcal{F} . Thus even if we let \mathcal{F} span its entire range the resultant \mathcal{D} does not span the entire space of observations, and this why a liberal attitude toward \mathcal{F} still yields interesting results.

To implement our liberal approach toward \mathcal{F} consider the class of probability distributions for \mathcal{F}

$$p_s(\mathcal{F})d\mathcal{F} = \alpha f_s(\alpha\mathcal{F})d\mathcal{F} . \quad (5)$$

Any normalizable distribution must fall off for large \mathcal{F} , but as $\alpha \rightarrow 0$ the region of significant probability density gets larger and larger and the variance of this prior distribution increases. However the likelihood function of eq. (3) will remain well behaved since for fixed \mathcal{D} , $p_2(\mathcal{D}|\mathcal{T} + \mathcal{F})$ will fall off. Note that \mathcal{T} and \mathcal{F} cannot exactly cancel since they have different spectra. Thus in the limit of large variance we may replace $p_s(\mathcal{F})$ with $p_s(0)$, and since this constant is independent of \mathcal{P} it will not enter into any likelihood ratio and we may drop it. Of course we require that $p_s(0) \neq 0$, however this is true of any reasonable distribution. For large variance we have

$$L(\mathcal{P}; \mathcal{D}, \mathcal{F}) \rightarrow L^*(\mathcal{P}; \mathcal{D}) \propto \int \int p_2(\mathcal{D}|\mathcal{T} + \mathcal{F}) p_1(\mathcal{T}|\mathcal{P}) d\mathcal{T} d\mathcal{F} . \quad (6)$$

Thus we see that this likelihood function is just what we would get for a uniform prior for \mathcal{F} . Such a uniform prior may include correlations between the different patches or correlations between different components of a multi-component foreground and the determinant of this correlation matrix may enter into $p_s(0)$, however this will not effect the likelihood ratio of eq. (6). Thus we are lead to our first interesting result:

- In the limit of large variance all prior distributions for the foreground contamination yield the same likelihood function.

We may simplify Eq. (6) still further if we make the assumption that the probability distribution of detector noise does not depend on the amplitude of the signal, i.e.

$$p_2(\mathcal{D}|\mathcal{T} + \mathcal{F}) = f_2(\mathcal{D} - \mathcal{T} - \mathcal{F}) . \quad (7)$$

To see how this helps let us denote \mathcal{D} 's dof which are independent of \mathcal{F} by \mathcal{D}^{ind} and the dependent dof by \mathcal{D}^{dep} . The MBR contribution to the signal may be similarly decomposed, but the foreground of course has no independent part. We may thus rewrite Eq. (7) as

$$p_2(\mathcal{D}|\mathcal{T} + \mathcal{F}) = f_2(\mathcal{D}^{\text{ind}} - \mathcal{T}^{\text{ind}}, \mathcal{D}^{\text{dep}} - \mathcal{T}^{\text{dep}} - \mathcal{F}) . \quad (8)$$

If we substitute this into eq. (6) and change one of the variables of integration from \mathcal{F} to \mathcal{D}^{dep} we find

$$L^*(\mathcal{P}; \mathcal{D}) \propto \int \left[\int p_2(\mathcal{D}|\mathcal{T} + \mathcal{F}) d\mathcal{D}^{\text{dep}} \right] p_1(\mathcal{T}|\mathcal{P}) d\mathcal{T} \propto L(\mathcal{P}; \mathcal{D}^{\text{ind}}) . \quad (9)$$

since the term in square brackets is just the marginal distribution of \mathcal{D}^{ind} obtained by integrating out \mathcal{D}^{dep} . This gives us our second interesting result

- The likelihood function constructed by using all of the data and a prior distribution for the foreground contamination with very large variance is equal to the likelihood function obtained by considering only that linear combination of the data which is independent of the foreground contamination.

One can obtain \mathcal{D}^{ind} by projecting out all linear combinations of the spectrum of the different foreground components. This last identity is of great practical use since it reduces the dimensionality of the data space and can make computations of the likelihood function much more tractable. Of course this idea of projecting out the foreground contamination is not a new one. For example, it has been used in constructing the “reduced galaxy” (RG) anisotropy map of the DMR experiment. However the above equality does give an added justification for constructing such RG data sets, as it shows that by doing so one is not throwing out any information except for ones assumptions about the properties of the foreground contamination.

The likelihood function L^* is obtained by integrating over arbitrarily large foreground contamination, which is certainly not true. Given some true distribution of foreground contamination, $p_s(\mathcal{F})$, under what circumstances will L^* give a good approximation to the true likelihood function? As mentioned before, p_z in Eq. (3) will regulate the integral over \mathcal{F} if p_s does not. Since p_z will start falling roughly when the foreground emission exceeds either the detector noise or the observed signal it follows that the condition for L^* to be an accurate representation of L is roughly that the variance in the foreground emission exceeds greatly either the detector noise or the observed signal.

In obtaining these results we have made no assumptions about the statistical properties of the MBR anisotropy or the foreground contamination, although one will have to make some assumptions about the former in order to compute L^* . In particular we haven’t assumed anything is Gaussian. We have assumed that the detector noise is independent of the amplitude of the signal, however we think this likely to be a very good approximation. More important assumptions were made about the experimental apparatus, in particular we have assumed that the effective window function is the same for all channels of a given patch. For some experiments, such as the DMR, this is an excellent approximation, while for other such as the ACME experiment it is not. The degree to which this varying window can confuse spatial and spectral dependencies will depend on the spatial distributions of the emission. One can alleviate this problem if the region of the sky is oversampled by the experiment, in which case one can bin the data into synthetic beam patterns which are the same for all channels. Finally we have assumed that we know the spectrum of the various foreground contaminants. Here it is important only that deviations from the assumed spectra are not sufficient that the \mathcal{D}^{ind} used in the likelihood function could be significantly contaminated by foreground

emission. Here we are making some assumption about the amplitude of foreground emission, but if our assumed spectrum is fairly accurate then this is a much less stringent constraint than assuming that the foreground emission is small compared with the observed signal. In the microwave region there is fairly small uncertainty in the spectra of free-free emission, however for synchrotron and dust emission this is not the case. Hopefully data from the DIRBE and FIRAS experiments on COBE will show a fairly universal spectra for dust emission at microwave frequencies. We do know that at longer wavelengths synchrotron emission exhibits an unfortunately broad range of spectral indices and this can lead to large uncertainty in constructing \mathcal{D}^{ind} .

Gaussian Statistics

Now we can apply this analysis specifically to the case where the MBR anisotropies are Gaussian random noise with some unknown parameters and the detector noise is also Gaussian. Let us represent the data, \mathcal{D} , by a set of numbers $\Delta_{(a,i)}$ which give the signal in channel a of patch i . We may similarly represent the contribution of the foreground contamination to each such patch and channel as Δ^f . Thus the MBR anisotropy plus detector noise, $\mathcal{T} + \mathcal{N}$, is given by $\Delta_{(a,i)} - \Delta_{(a,i)}^f$, which has the correlation matrix

$$\langle (\Delta_{(a,i)} - \Delta_{(a,i)}^f)(\Delta_{(b,j)} - \Delta_{(b,j)}^f) \rangle = C_{(a,i)(b,j)} \quad C_{(a,i)(b,j)} = C_{ij}^{\text{mbr}} + \sigma_{(a,i)}^2 \delta_{ij} \delta_{ab} \quad (10)$$

where $\sigma_{(a,i)}$ gives the instrumental noise and C_{ij}^{mbr} gives the expected correlation in MBR fluctuations. We will assume that we have some model which determines C_{ij}^{mbr} modulo the value of some parameters which are the \mathcal{P} of the previous discussion. The MBR anisotropy contribution to the signal and the instrumental noise are assumed Gaussianly distributed with zero mean and are thus fully determined by their respective correlation matrices, C_{ij}^{mbr} and $\sigma_{(a,i)}^2 \delta_{ij} \delta_{ab}$. The foreground emission may be written as a sum over the different foreground emission processes (e.g. free-free, synchrotron, dust), i.e.

$$\Delta_{(a,i)}^f = \sum_{\alpha=1}^{N_f} A_{[\alpha,i]} F_{\{\alpha,a\}} \quad (11)$$

where α labels the process, $A_{[\alpha,i]}$ gives the amplitude of emission of process α in beam i , and $F_{\{\alpha,a\}}$ gives the contribution per unit amplitude of process α in channel a . The quantities $F_{\{\alpha,a\}}$ are assumed known, however the amplitudes $A_{[\alpha,i]}$ are not. It is these $A_{[\alpha,i]}$'s which represent the \mathcal{F} in the previous subsection.

One can explicitly compute the components of $\Delta_{(a,i)}$ which are independent of the foreground by solving the system of equations

$$\sum_{a=1}^{N_{\text{ch}}} z_a F_{\{\alpha,a\}} = 0 \quad \alpha = 1, \dots, N_f, \quad (12)$$

say with a general purpose eigensystem solver. As long as one is able to distinguish the different components of foreground contamination, i.e. as long as $F_{\{a,i\}}$ are linearly independent, there will be an $(N_{\text{ch}} - N_f)$ -dimensional space of solutions of eq. (12). If one chooses a basis for this space, $z_a^{(r)}$, where r labels the basis vectors, then one obtains a set of coordinates on the data subspace which are independent of the foreground:

$$\Delta_{[r,i]}^{\text{ind}} = \sum_{a=1}^{N_{\text{ch}}} z_a^{(r)} \Delta_{\{a,i\}}^{\text{ind}} . \quad (13)$$

The correlation function on \mathcal{D}^{ind} is

$$\langle \Delta_{[r,i]}^{\text{ind}} \Delta_{[s,j]}^{\text{ind}} \rangle = C_{[r,i][s,j]}^{\text{ind}} = \sum_{a=1}^{N_{\text{ch}}} \sum_{b=1}^{N_{\text{ch}}} z_a^{(r)} C_{\{a,i\}\{b,j\}} z_b^{(s)} \quad (14)$$

so the likelihood functions of eq. (9) are

$$L^*(\mathcal{P}; \mathcal{D}) \propto L(\mathcal{P}; \mathcal{D}^{\text{ind}}) \propto \frac{\exp\left(-\frac{1}{2} \chi_{\text{ind}}^2\right)}{\|C_{[r,i][s,j]}^{\text{ind}}\|} \quad (15)$$

where $\| \cdot \|$ indicates the determinant and

$$\chi_{\text{ind}}^2 = \sum_{i=1}^{N_p} \sum_{j=1}^{N_p} \sum_{r=1}^{N_{\text{ch}}-N_f} \sum_{s=1}^{N_{\text{ch}}-N_f} \Delta_{[r,i]}^{\text{ind}} C_{[r,i][s,j]}^{\text{ind}^{-1}} \Delta_{[s,j]}^{\text{ind}} . \quad (16)$$

Of course, χ_{ind}^2 is just the chi-square statistic which for the correct choice of $C_{[r,i][s,j]}^{\text{ind}}$ will be distributed like a χ^2 -distribution with $N_p \times (N_{\text{ch}} - N_f)$ dof.

3. Example: South Pole 91

In this section, we will flesh out the formalism we've set up with an example. We will restrict our analysis to testing one particular theory, cold dark matter (CDM) with a Harrison- Zel'dovich initial spectrum. We'll focus on an experiment that's been analyzed a number of times already, so at least the first part of our discussion, which assumes there are no foreground sources, should be familiar to many. This experiment, which we refer to as SP91, was performed at the South Pole in 1991 with the ACME-HEMT telescope, and the results were published in Gaier *et al.* (1992).

One Patch; One Frequency Channel

We start by considering the simplest possibility: An experiment measures the temperature difference in one region of the sky at one frequency. This one measurement, call it Δ , is thus the

full set of our data \mathcal{D} . The first thing we need to know is what is the temperature distribution – the full set \mathcal{T} is now simply one number T – predicted by the theory, $p_1(T|Q_{\text{rms-ps}})$. The one parameter in the theory [the set previously denoted \mathcal{P}] is the normalization, which we'll take as $Q_{\text{rms-ps}}$, the average quadrupole over an ensemble of universes. Since perturbations are Gaussian in standard cold dark matter, we can write

$$p_1(T|Q_{\text{rms-ps}}) = \frac{1}{\sqrt{2\pi C^{\text{mbr}}}} \exp \left\{ -\frac{1}{2} \frac{T^2}{C^{\text{mbr}}} \right\}. \quad (17)$$

The width of the Gaussian here, C^{mbr} , depends on both the theory and the experimental configuration of the beam. In particular,

$$C^{\text{mbr}} = \sum_{l=2}^{\infty} \frac{2l+1}{4\pi} C_l W_l \quad (18)$$

where C_l is the coefficient of the Legendre polynomial $P_l(\hat{n}_1 \cdot \hat{n}_2)$ when $C(\hat{n}_1 \cdot \hat{n}_2) \equiv \langle T(\hat{n}_1)T(\hat{n}_2) \rangle$ is expanded in a series of such polynomials. That is, the C_l 's are given by the theory. The theory's one free parameter, $Q_{\text{rms-ps}}$, is related to C_2 via: $C_2 = 4\pi Q^2/5$. Meanwhile, the window function W_l is solely determined by the experimental beam size and chopping strategy. For the SP91 experiment (Bond *et al.* 1991)

$$W_l = \exp \left\{ -l(l+1)\theta_s^2 \right\} \frac{16\pi}{2l+1} \sum_{m=-l}^l H_0^2(m\phi_A) Y_{lm}^2(\theta_z, 0) \quad (19)$$

where $\theta_z = 27.75^\circ$; $\phi_A = 1.5^\circ / \sin(\theta_z)$; $\theta_s = 0.425 \times 1.35^\circ$ [for the highest frequency channel we are discussing at present]; and H_0 is the Struve function of order zero. The C_l 's for CDM and W_l for SP91 are plotted in Figure 1. Once the C_l and W_l are given, it is straightforward to combine them and compute the expected variance C^{mbr} . For standard CDM and SP91, we find $C^{\text{mbr}} = (43Q_{\text{rms-ps}}/17)^2$. [Recall that COBE-normalized CDM has $Q_{\text{rms-ps}} = 17\mu\text{K}$.]

The next step is to account for experimental errors by calculating $p_2(\Delta|T)$. We assume the errors are Gaussian, so that

$$p_2(\Delta|T) = \frac{1}{\sqrt{2\pi} \sigma_{\text{exp}}} \exp \left\{ -\frac{1}{2} \frac{(\Delta - T)^2}{\sigma_{\text{exp}}^2} \right\}. \quad (20)$$

Here, σ_{exp} is the variance of measurements in the 'lab' where there are no other sources [cosmic or otherwise] contributing to the signal. So, eq. (20) simply tells us that if σ_{exp} were very small, the observed value Δ would be very close to the actual value on the sky, T . On the other hand, if the noise is significant, the observed value could differ significantly from the sky value. One of




FIG. 1. The window function W_l for the SP91 experiment and the C_l 's in CDM with $h = 0.5$; $\Omega_B = 0.06$; $n = 1$. Plotted is $l(l+1)C_l/6$ which is equal to one if only the Sachs-Wolfe effect is considered. The dashed line is an approximation to the SP91 filter function [assuming square well chopping] which is seen to be off by as much as 30%. Also shown is the filter function for the MAX experiments to be analyzed in the next section.

the exciting things about the SP91 experiment is that σ_{exp} was of order $20 - 30\mu\text{K}$, significantly below the signal predicted by CDM.

Now that p_1 and p_2 are given, we can convolve the two as required by eq. (18) to form $p_3(\Delta|Q_{\text{rms-ps}})$. The integral over T is readily performed and we find

$$p_3(\Delta|Q_{\text{rms-ps}}) = \frac{1}{(2\pi)^{N/2}} \|C\|^{-1/2} \exp \left\{ -\frac{1}{2} \Delta C^{-1} \Delta \right\} \quad (21)$$

where N is the number of measurements [here just one] and the correlation “matrix” [in this case one by one]

$$C \equiv \langle |T + \mathcal{N}|^2 \rangle = C^{\text{mbr}} + \sigma_{\text{exp}}^2. \quad (22)$$

Figure 2 shows p_3 as a function of the observed temperature Δ for several values of the normal-

ization $Q_{\text{rms-ps}}$. Here is a good time to hearken back to our discussion following eq. (1). There we argued that there were two classes of ways one might constrain the parameters in a theory. Let us illustrate these two classes with the aid of Figure 2. First, we could require $Q_{\text{rms-ps}}$ to be such that the probability density of the observed value of Δ is reasonably high. For example, if Δ were $0 \mu\text{K}$, we might allow the parameters $Q_{\text{rms-ps}} = 10\mu\text{K}$ and $Q_{\text{rms-ps}} = 20\mu\text{K}$, but frown on $Q_{\text{rms-ps}} = 30\mu\text{K}$, because the probability density of $\Delta = 0 \mu\text{K}$ is unacceptably low. If, however, Δ was observed to be $200\mu\text{K}$ we would rule out all three values of $Q_{\text{rms-ps}}$, because the probability density is too low in each case. Now consider the second method: the likelihood approach. In this approach, we compare the probability density of Δ for different values of $Q_{\text{rms-ps}}$ and throw out values of $Q_{\text{rms-ps}}$ which have likelihoods significantly smaller than the “best” values of $Q_{\text{rms-ps}}$. In our artificial example here, with only three values of $Q_{\text{rms-ps}}$, this means that if Δ were $200\mu\text{K}$, we would throw out $Q_{\text{rms-ps}} = 10, 20\mu\text{K}$, but keep $Q_{\text{rms-ps}} = 30\mu\text{K}$ since this is the value of $Q_{\text{rms-ps}}$ at which the likelihood [for the observed Δ] is maximum. There is something unsatisfactory about this: We are accepting a value of the parameter which gives the best fit, but it is not a particularly good fit. So although we will follow the second approach to find the best fit $Q_{\text{rms-ps}}$, we will use the first approach to “check” the goodness of this best fit.

Many Patches; One Frequency Channel

The nine point scan of SP91 reported temperature differences for nine patches on the sky. The discussion above is easily generalized to the multi-patch case. The data \mathcal{D} which before was a single measurement Δ is now a series of measurements, $\Delta_i, i = 1, \dots, N_{\text{patch}} = 9$. We have seen that all we need to calculate p_3 [or the likelihood function] is the correlation matrix, C , which now is 9×9 . The experimental errors are assumed independent so

$$C_{ij} = \langle T_i T_j \rangle + \delta_{i,j} \sigma_{\text{exp},i}^2. \quad (23)$$

The off-diagonal elements of the theoretical correlation matrix $\langle T_i T_j \rangle$ are given by the same sum in eq. (18), with

$$W_{l,i,j} = \exp \{ -l(l+1)\theta_z^2 \} \frac{16\pi}{2l+1} \sum_{m=-l}^l \cos(m(\phi_i - \phi_j)) H_0^2(m\phi_A) Y_{lm}^2(\theta_z, 0) \quad (24)$$

where ϕ_i is the azimuthal angle corresponding to the center of the i^{th} patch [the polar angle θ_z is kept constant throughout the scan].

There is one further complication that must be dealt with before we can show the likelihood function. The experimenters subtract from each scan an “offset” and a “drift.” That is, they




FIG. 2. The probability distribution for the observed temperature in a single patch of the SP91 experiment, $p_3(\Delta|Q_{\text{rms-ps}})$, for three different values of the CDM normalization, $Q_{\text{rms-ps}}$. The distribution, which is the same for $-\Delta$, is normalized so its integral over all Δ is one.

subtract from each scan the best fit line. So the temperature differences they are actually reporting are

$$\Delta'_i \equiv \Delta_i - (m\phi_i + b). \quad (25)$$

There are several ways of accounting for this. Bond *et al.* (1991) assumed a uniform prior for the average and gradient and integrated them out. Another possibility is to note that the 8th and 9th patch are really redundant, since they are fixed by the requirement that the mean and gradient vanish. Thus we could simply project onto the seven-dimensional space of measurements. All of this has a familiar ring to it: these are precisely the two alternatives we found to be equivalent when we talked about foreground. So as a warmup exercise to subtracting off foreground, let's apply the formalism we set up in section 2 to subtract off the mean [we won't worry about the gradient in this discussion, although we do subtract it off to get our final results]. Analogous to

eq. (12) we want the final temperature set to have zero mean. If all the errors were the same, this would mean $\sum_{i=1}^{N_p} z_i = 0$. Putting in the weighting factors leads to

$$\sum_{i=1}^{N_p} \frac{z_i}{\sigma_{\text{exp},i}^2} = 0. \quad (26)$$

This equation is satisfied for eight linearly independent $z_i^{(r)}$, $r = 1, \dots, 8$. Once we find these eight vectors, we can then form the mean-subtracted temperatures

$$\Delta'_r = \sum_{i=1}^9 z_i^{(r)} \Delta_i. \quad (27)$$

To calculate p_3 , we can still use eq. (21), but now the correlation matrix that enters is the reduced correlation matrix

$$C_{rs} \equiv \langle \Delta'_r \Delta'_s \rangle = \sum_{i,j=1}^{N_p} z_i^{(r)} z_j^{(s)} C_{ij}. \quad (28)$$

If we assume a prior distribution function, $p_4(Q_{\text{rms-ps}})$, we can convert p_3 into a posterior distribution for $Q_{\text{rms-ps}}$, $p_5(Q_{\text{rms-ps}}|\mathcal{D})$, using eq. (2). (We see from eq. (2) that the functional dependence of the likelihood function on the model parameter is the same as the posterior probability density for the parameters in a Bayesian analysis if the assumed prior is uniform in the parameters. Since others (e.g. Srednicki *et al.* 1993) have put a Bayesian interpretation on their analysis we will find it convenient to discuss various prior distributions, and use the resultant posterior probability densities to set limits on $Q_{\text{rms-ps}}$. However one should remember that when we consider a uniform prior, we might just as well be doing a likelihood analysis as a Bayesian analysis. The names are changed but they are mathematically equivalent.) Figure 3 shows this posterior distribution function for the highest frequency (4th) channel of the SP91 experiment, in the case of three different priors. For the 4th channel SP91 data, no matter which prior is used, p_5 falls off fairly quickly, which led a number of groups to place fairly stringent upper limits. The exact limit does depend sensitively on the prior though. If we define an upper limit via

$$\int_0^{Q_{\text{rms-ps}}^{\text{upper limit}}} dQ p_5(Q) = 0.95, \quad (29)$$

then the upper limits are $Q_{\text{rms-ps}}^{\text{upper limit}} = 9, 14, 20 \mu\text{K}$ for $p_4(Q_{\text{rms-ps}}) = 1/Q_{\text{rms-ps}}, 1, Q_{\text{rms-ps}}$, respectively. One way to assess the upper limit associated with a given prior is to calculate the *level of significance* of the test. This tells us how often we would rule out $Q_{\text{rms-ps}}^{\text{upper limit}}$ if that was the true value. For example, when $p_4(Q_{\text{rms-ps}}) = 1$ leading to an upper limit of $Q_{\text{rms-ps}}^{\text{upper limit}} = 14 \mu\text{K}$,




FIG. 3. The posterior distribution for $Q_{\text{rms-ps}}$ using the highest frequency channel of the SP91 experiment. The three different curves correspond to three different choices of priors: $p_4 = 1/Q_{\text{rms-ps}}$, 1 , $Q_{\text{rms-ps}}$. The first such prior has been cut off at $Q_{\text{rms-ps}} = 2\mu\text{K}$ to make it normalizable.

we can ask: If the Universe truly had $Q_{\text{rms-ps}} = 14\mu\text{K}$, and many different experiments were done, how often would someone using this test rule out $Q_{\text{rms-ps}} = 14\mu\text{K}$? The levels of significance of the tests with the priors $p_4 = 1/Q_{\text{rms-ps}}$, 1 , $Q_{\text{rms-ps}}$ are .11, .02, .002, respectively. Thus a prior which favors high values of $Q_{\text{rms-ps}}$ leads to very low levels of significance; i.e. it leads to upper limits which are too stringent. For this reason, $p_4 = 1$ is often chosen, and we will stick to this choice for the rest of our discussion.

Many Patches; Many Frequency Channels

We now account for the fact that SP91, like many modern anisotropy experiments, took measurements at a number of different frequency channels. In the SP91 experiments there were four frequency channels, spanning the range 25 – 35 GHz. Thus, the data set \mathcal{D} now consists of $\Delta_{(a,i)}$, $a = 1, \dots, 4$; $i = 1, \dots, 7$ [recall that the mean and gradient are subtracted out of each

channel]. To construct the posterior distribution p_5 [from now on we'll use $p_4 = 1$], we must therefore form the correlation matrix

$$C_{(a,i)(b,j)} = \langle T_{a,i} T_{b,j} \rangle + \delta_{ij} \delta_{ab} \sigma_{\text{exp},(a,i)}^2. \quad (30)$$

Ordinarily, the expected cosmic signal would be the same in each channel, since the temperature differences are frequency independent. In SP91, there is a small complication owing to the different widths of the beams in the different channels. We have accounted for this by allowing θ_s in the filter function of eq. (19) to be channel dependent.

Figure 4 shows the posterior distribution for $Q_{\text{rms-ps}}$ given the 4-channel data in the SP91 experiment. Taken at face value, this data clearly seems to indicate a detection, i.e. $Q_{\text{rms-ps}} = 0$ is ruled out. Applying the test in eq. (29) to determine an upper limit and a similar one to determine the lower limit, i.e.

$$\int_{Q_{\text{rms-ps}}^{\text{lower limit}}}^{\infty} dQ p_5(Q) = 0.05, \quad (31)$$

we find

$$Q_{\text{rms-ps}} = 10_{-4}^{+11} \mu\text{K} \quad (32)$$

at the 95% confidence level. Eq. (32) tells us that the best fit for this experiment is at about $Q_{\text{rms-ps}} = 10 \mu\text{K}$. Is this “best fit” a good fit? When there was only one data point we could simply look at the distribution p_3 and see whether or not the distribution function was acceptably high at the actual values observed. Unfortunately, it is harder to do this in a 28 dimensional space. Instead, a good number to look at for these purposes is the χ^2 , defined in eq. (16). The observed χ^2 should be of order the number of degrees of freedom, in SP91 (4 channels) \times (7 patches) – (1 normalization parameter) = 27, with a standard deviation of $\sqrt{2 \times N_{\text{dof}}} = 7$. However, in this case the observed $\chi^2 = 46$ for the four-channel data. This tells us that our best fit is not a very good fit at all, and we better look at a different theory or at least at the possibility of other sources. We turn next to this latter possibility.

One Source of Foreground: Marginalization

We now allow the possibility of foreground sources. The observed signal then is the sum of detector noise, foreground sources, and cosmic signal. As in Eq. (3), the distribution function for the data is

$$p_3(\mathcal{D}|Q_{\text{rms-ps}}) = \frac{1}{(2\pi)^{(N/2)}} \|C\|^{-1/2} \int d\mathcal{F} p_6(\mathcal{F}) \exp \left\{ -\frac{1}{2} (\mathcal{D} - \mathcal{F}) C^{-1} (\mathcal{D} - \mathcal{F}) \right\}. \quad (33)$$




FIG. 4. The posterior distribution for $Q_{\text{rms-ps}}$ from the four channel, nine patch scan of SP91. The solid line assumes no foreground contribution. The dashed line assumes a best fit free-free emission with frequency dependence as in eq. (34). Dotted line shows the result if free-free emission is subtracted off. Dot-dashed line shows the results if synchrotron and free-free are subtracted off.

The number of independent measurements $N = 28$, one for each patch and channel. Until now we have been implicitly assuming a delta function for the foreground prior: $p_6 \propto \delta(\mathcal{F})$, no foreground. In principle, \mathcal{F} is a set of 28 numbers, one for each patch and channel. Let us first suppose that there is only one component to the foreground, free-free emission, say. Then, we expect the signal in each channel to scale as

$$F_a \propto \nu_a^{-2.1} \quad (34)$$

where ν_a is the frequency associated with the a^{th} channel. The distribution function now becomes

$$p_3(\mathcal{D}|Q_{\text{rms-ps}}) = \frac{1}{(2\pi)^{(N_{\text{ch}}N_p/2)}} ||C||^{-1/2} \int \prod_{i=1}^{N_p} (dA_i) \times p_6(\{A_i\}) \exp \left\{ -\frac{1}{2} (\Delta_{(a,i)} - A_i F_a) C_{(a,i)(b,j)}^{-1} (\Delta_{(b,j)} - A_j F_b) \right\}. \quad (35)$$

If the foreground contribution is known in one of the channels, eq. (34) allows us to determine the contribution in all channels. Thus, instead of 28 free parameters, we have added only seven. Other maps of the region, radio maps or infrared maps, might give a reasonable prior p_6 , so that the integral in eq. (35) could be done and translated into a distribution for $Q_{\text{rms-ps}}$. In the absence of these, we will use a uniform prior for the A_i , which we've argued is equivalent to considering only those linear combinations of the temperatures which are independent of the free-free emission. Before we do this, though, it is interesting to consider another possible approach to the integral in eq. (35). Besides the prior, the only dependence on the free-free emission comes in the quadratic polynomial in the exponential. One could easily minimize the argument of the exponential, thereby maximizing the distribution. If the integrand is sharply peaked, then this "best fit foreground" should be a good approximation to the distribution. Figure 4 shows the likelihood function obtained in this way. The best fit free-free leaves little room for anything else! The lesson to be learned from this example is *not* that a very stringent upper limit has been set. Rather, it is important to note the difference that foreground sources can make. The likelihood for $Q_{\text{rms-ps}}$ with best fit foreground differs significantly from the likelihood with no foreground. Since the foreground can make a world of difference and since we really don't have much prior information about the foreground, one reasonable approach is to keep only the parts of the signal that are independent of foreground, i.e. to marginalize.

We subtract off components dependent on foreground by solving eq. (12); we'll do this first assuming a foreground spectrum as in eq. (34). The four frequency channels in the SP91 experiment are centered at $\nu = 26.25, 28.75, 31.25, 33.75$ GHz. Therefore, we must solve

$$\sum_{a=1}^4 z_a \left(\frac{26.25}{26.25 + (a-1)2.5} \right)^{2.1} = 0. \quad (36)$$

There are three independent solutions to this equation. So the four temperatures in the four channels have been reduced to three temperatures –linear combinations of the four channels:

$$\begin{aligned} \Delta_{(1,i)}^{\text{ind}} &= -2.5\Delta_{(1,i)} + 1.3\Delta_{(2,i)} + 1.2\Delta_{(3,i)} + 1.0\Delta_{(4,i)} \\ \Delta_{(2,i)}^{\text{ind}} &= 0.0\Delta_{(1,i)} - 6.7\Delta_{(2,i)} + 9.6\Delta_{(3,i)} - 1.9\Delta_{(4,i)} \\ \Delta_{(3,i)}^{\text{ind}} &= 0.0\Delta_{(1,i)} - 2.2\Delta_{(2,i)} - 0.7\Delta_{(3,i)} + 3.9\Delta_{(4,i)}. \end{aligned} \quad (37)$$

The index i here still labels patches, as we perform this transformation in each patch. The normalization is arbitrary; we have chosen it so that for MBR, with its flat spectrum, $\Delta^{\text{ind}} = \Delta$. From here on in, the construction of the likelihood function mimics the process we went through with mean-subtraction. The main step is to construct the reduced correlation matrix as in eq. (14).

Figure 4 shows the likelihood function thus obtained. Perhaps the most surprising feature of the marginalized likelihood function is that it leads to a very weak upper limit:

$$Q_{\text{rms-ps}}^{\text{upper limit}} = 52 \mu\text{K} . \quad (38)$$

This is a bit surprising because one might have reasoned as follows: Perhaps the lower limit coming from the four channel data with no foreground should be ignored because part of the signal could have come from something else. But at least the upper limit should be reliable even if foreground is accounted for. We see now that this reasoning is wrong! Not only can foreground confuse us into thinking there is a cosmic signal when in reality none exists, but also foreground can take away part of a cosmic signal. Foreground can be negative when the cosmic signal is positive, so by ignoring foreground, one can *underestimate* the cosmic signal, and it is this possibility that leads to the weak upper limit. Some may find the weak upper limit of eq. (38) overly conservative since they would consider the possibility of the foreground emission decreasing the variance in channel 4 to be negligible. We will argue against this point of view in below in §5.

Several Sources of Foreground

Until now we have allowed only one source of foreground. What happens if we allow several sources? The answer is disquieting. We now show that with more than one source of foreground and no prior information about the amplitudes, no interesting limits can be placed on the parameter $Q_{\text{rms-ps}}$.

If we subtract a synchrotron component [assuming $F \propto \nu^{-2.7}$] in addition to the free-free component, the two independent temperatures in each patch are

$$\begin{aligned}\Delta_{(1,i)}^{\text{ind}} &= 11.3\Delta_{(1,i)} - 21.8\Delta_{(2,i)} - 0.8\Delta_{(3,i)} + 12.3\Delta_{(4,i)} \\ \Delta_{(2,i)}^{\text{ind}} &= 0.0\Delta_{(1,i)} + 21.5\Delta_{(2,i)} - 54.7\Delta_{(3,i)} + 34.2\Delta_{(4,i)}.\end{aligned}\tag{39}$$

If three components – free-free, synchrotron, and dust [$F \propto \nu^{1.6}$] – then the one independent component in each patch is

$$\Delta_{(1,i)}^{\text{ind}} = 116\Delta_{(1,i)} - 420\Delta_{(2,i)} + 494\Delta_{(3,i)} - 189\Delta_{(4,i)}.\tag{40}$$

Recall that Δ^{ind} is the sum of the cosmic component (T^{ind}) and the detector noise (N^{ind}). As we've said, the cosmic component contributes equally in each channel so in a given patch $T^{\text{ind}} = T$. However, the detector noise is completely uncorrelated in the different channels, so the large coefficients in Eqs. (39) and (40) tells us that the detector contribution to Δ^{ind} will be enormous. The cosmic signal will be dwarfed by the detector noise. As an example of this, Fig. 4 shows the likelihood function when both synchrotron and free-free are subtracted off. The likelihood function is extremely flat, and the upper limit is not very useful: $Q_{\text{rms-ps}}^{\text{upper limit}} = 200\mu\text{K}$.

These conclusions are not particular to SP91. For all the experiments we will analyze, subtracting off more than one source of contamination results in data swamped by noise and therefore useless for setting limits. This may be due to the small number of channels [three for MAX-ACME and four for the South Pole scans]; perhaps more channels would improve the situation. And of course we have not used any information from other maps which might further constrain the sources.

Perhaps subtracting off one source is enough though. One of the main incentives for considering foreground when analyzing the SP91 results was the large χ^2 for the best fit $Q_{\text{rms-ps}}$ using the four channel data [$\chi^2 = 46$ for 27 dof]. When only one source of foreground is included $\chi^2 = 28$ for 20 dof, a significantly better fit. In all of the experiments we analyze here, subtracting off one foreground component makes the best fit $Q_{\text{rms-ps}}$ a better fit.

4. Other Experiments

In this section we apply the method developed and explained in the previous two sections to three other recent experiments. Along with the nine point scan of the South Pole experiment, there was a 13– point scan (Schuster, *et al.* 1992) which probed a nearby region of the sky. At larger frequencies, the Millimeter Anisotropy Experiment (MAX) recently reported results of two scans: one around the region of μ -Pegasus and the other near the star Gamma Ursa Minoris (GUM).

Experiment	Foreground	$Q_{\text{rms}}/\mu\text{K}$	$\chi^2/\text{d.o.f.}$	$P(\chi^2)$
SP 9-point scan	None	10^{+11}_{-4}	46/27	.012
	Free-Free	22^{+25}_{-17}	28/20	.12
	Free-Free + Synchrotron	< 184	18/13	.16
SP 13-point scan	None	12^{+8}_{-4}	58/43	.06
	Free-Free	14^{+16}_{-12}	40/32	.17
MAX: μ -Peg scan	Warm Dust	8^{+9}_{-4}	35/41	.73
	Warm Dust + Free-Free	< 86	19/20	.52
MAX: GUM scan	None	29^{+14}_{-7}	149/113	.02
	Cold Dust	29^{+16}_{-9}	98/75	.05
All Four Experiments	One Component	18^{+8}_{-5}	208/168	.02

Table 1. Results of Bayesian analyses of four experiments. $Q_{\text{rms-ps}}$ lists the best fit $Q_{\text{rms-ps}}$ with the error bars indicating the 95% upper and lower limits using Bayesian analysis with a uniform prior for $Q_{\text{rms-ps}}$. The column headed χ^2/dof lists the χ^2 of the best fit $Q_{\text{rms-ps}}$. The last column tabulates the probability of getting a χ^2 as large as or larger than this for the stated number of degrees of freedom.

Figure 5 shows the results of calculating the likelihood function for these experiments with and without marginalization. Table 1 gives quantitative upper and lower limits and presents the χ^2 for the best fit value of $Q_{\text{rms-ps}}$. Also tabulated is the probability of getting a χ^2/dof this large, which we'll take as a measure of goodness-of-fit.

The raw Schuster data indicates a detection, but as Table 1 shows, the best fit for this data is quite poor if we assume it is all cosmic background. Specifically, the probability of getting a $\chi^2/\text{dof} = 58/43$ or larger is only .06. The situation improves somewhat if one foreground component is subtracted off. In this case, though, the detection becomes less significant: the likelihood function




FIG. 5. The likelihood functions for the SP 13 point scan and two MAX scans. The solid line labeled μ -Peg is the likelihood for μ -Peg when warm dust is marginalized; the bottom-most dashed line is for μ -Peg when both warm dust and free-free [with $T \propto \nu^{-2.1}$] are marginalized. The solid line labeled SP13 is for the raw data presented in the thirteen point scan of the South Pole experiment; the dashed line which is broader but peaks in roughly the same place is for the same experiment with free-free marginalized. The solid curve for the GUM scan is the likelihood using the raw data; the nearby dashed curve shows the likelihood when cold dust is marginalized. The heavy solid line is the product of the likelihood functions for the four experiments, each with one foreground component marginalized.

at $Q_{\text{rms-ps}} = 0$ is still 60% of its maximum at $Q_{\text{rms-ps}} = 12\mu\text{K}$.

The MAX experiment has a slightly different chopping procedure than the South Pole experiments, so Eq. (24) is replaced by

$$W_{l,ij} = \left(\frac{3.34}{2}\right)^2 \exp\{-l(l+1)\theta_s^2\} \frac{16\pi}{2l+1} \sum_{m=-l}^l \left(\frac{\sin(m\phi_{*,i}/2)}{m\phi_{*,i}/2}\right) \left(\frac{\sin(m\phi_{*,j}/2)}{m\phi_{*,j}/2}\right) \\ \times \cos(m(\phi_i - \phi_j)) J_1(m\phi_{A,i}) J_1(m\phi_{A,j}) Y_{lm}^*(\theta_{z,i}, 0) Y_{lm}(\theta_{z,j}, 0). \quad (41)$$

Here J_1 is the Bessel function of order one. For both scans the beam smearing angle $\theta_s = 0.425 \times 0.5^\circ$, and the chop amplitude was $\phi_A \sin(\theta_z) = 0.65^\circ$. The μ -Peg scan was taken at constant $\theta_z = 65.45^\circ$, with an angle of $\phi_* = 0.285^\circ / \sin(\theta_z)$ between the center of each spatial patch. The GUM scan was binned so that $\phi_* = 1.125^\circ$. The scan around GUM took data at several strips; that is, θ_z was not constant. We account for this by allowing each patch to have $\theta_{z,i}$, a polar angle which varies from row to row [there are four rows in all]. Eq. (41) is based on the assumption that the scan took place at constant θ_z , whereas the scans really took place in a “bowtie” pattern. However, if the secondary chop was rapid enough, the constant azimuth approximation should be a good one. And indeed our results for the raw data alone agree with other preliminary analyses (Bond 1993; Srednicki, et al. 1993), the latter of which did account for the bowtie pattern. The prefactor in Eq. (41) is explained in Srednicki *et al.* (1993); it normalizes the signal so that the MAX filter really would report $T_1 - T_2$ if it scanned across a region with two different temperatures.

The μ -Peg data correlates very well with dust from the IRAS catalogue. The MAX team used two methods to extract information about the MBR from this data. First, the IRAS dust was directly subtracted, and residuals analyzed as MBR; second, the full data was simultaneously fit for dust and MBR. Subsequent analyses have used the residual data set. Both of these methods give similar answers. It is worth pointing out though that marginalizing with respect to dust [as we do in Figure 5] is yet another method: instead of subtracting off a best fit dust component [as the simultaneous fit does] marginalizing integrates over all possible dust contributions, weighting each by the internal data. This third method, represented by the solid line in Figure 5, gives results in striking agreement with the other two. The dotted line in Figure 5 [the most horizontal one] shows what happens if two components are subtracted off. Again, little information can be gleaned; the upper limit is $86\mu\text{K}$.

The likelihood for the GUM scan is also shown in Figure 5. The solid line is the likelihood for the raw data, which leads to a high normalization [recall that COBE’s normalization is now $Q_{\text{rms-ps}} = 17 \pm 3\mu\text{K}$]. Table 1 shows that this best fit is not a particularly good fit; however marginalizing with respect to cold dust leads to little change in either the best fit or the goodness of fit.

Finally the heavy solid line in Figure 5 shows the combined likelihood function of all four experiments we have analyzed. Here we have simply multiplied the four likelihoods together, in each case choosing the one-foreground-component-subtracted data. The normalization agrees eerily well with that of COBE, but this is not necessarily the number to focus on. Instead, we note that the likelihood function thus obtained is quite narrow; Table 1 shows that $Q_{\text{rms-ps}} = 18^{+8}_{-5}\mu\text{K}$, so the error bars are quite small. This is encouraging, because a valid criticism of this work would

be that we have been too liberal in our assumption about the amplitude of the foreground. The small error bars show that we have not thrown out too much information. This best fit is also a poor fit: the probability of getting a $\chi^2/\text{dof} = 208/168$ or larger is only 0.02 reflecting the fact that all of the individual experiments except for the μ -Peg scan exhibit relatively poor fits with large χ^2 's. The generally poor fit might be simply telling us that the theory we have chosen is not the correct one. A theory with more power on small scales would help fit the GUM results better. Nonetheless, the fact that the two MAX scans sample the anisotropy at the same angular scales and find very different results suggests that it will not be easy to fit the data by simply adjusting the angular power spectrum. The probable interpretation is that foreground has still not been adequately removed from the data. Perhaps the foreground has a very different spectrum than has been assumed here, say self-absorbed synchrotron or very cold dust. Of course we may have assumed not only the wrong power spectrum for the MBR but also the wrong statistical distribution. If the primordial anisotropy field is not Gaussian then one might expect a greater probability of having very different levels of anisotropies in one part of the sky than in another. This might explain the difference between the μ -Peg and the GUM results. Here it should be remembered that our analysis of the different experiments includes the uncertainty due to finite sampling. One cannot ascribe the apparent inconsistency between GUM and μ -Peg under the CDM model to finite sampling.

5. Are Our Limits Too Conservative?

In the preceding discussion we have taken into account the possibility of foreground emission by throwing away those components of measurements which might be effected by the sources of foreground we are considering. This process of marginalization makes no attempt whatsoever to determine the amount of foreground contamination from the microwave data themselves. As illustrated by both fig. 3 and the table, this procedure can indeed lead to very weak limits on the parameters characterizing the primordial anisotropy. Is throwing out all of this data really justified?

Clearly the limits produced by marginalization are conservative ones, as it allows for foreground contamination with an amplitude so large that it is unlikely to have produced the data that is being analyzed. One could reasonably try to limit the foreground contamination by insisting on goodness-of-fit to the microwave data and then using the constraints on the foreground to construct a prior for the foreground emission, i.e. p_8 of eq. (3). The marginalization procedure is simpler and more straightforward in that no spatial correlation for the foreground emission has to be constructed. In many cases marginalization will not lead to much less stringent limits than

those obtained by the Bayesian approach just described. The utility of marginalization is best considered in light of the commonly used alternative procedures:

- 1) analyze ignoring the possibility of foreground contamination,
- 2) analyze ignoring foreground after removing the best fit foreground, or
- 3) analyze after culling of data points which appear to be contaminated.

All of these procedures we would consider to be dangerous as we will now explain.

A good guiding principle in dealing with foreground is that one treat the detector noise and “foreground noise” (i.e. contamination) on an equal footing. If we have no good information about the foreground contamination from radio or IR measurements the only real differences between the two is that one has a good idea of the characteristics of the former and very little idea of the characteristics of the latter. In the MBR community it has been generally accepted that one should not use data whose amplitude is less than the sensitivity (i.e. noise level) to put limits on the MBR anisotropy below the sensitivity of the detector. We would like to generalize this to say that one shouldn’t set upper limits on primordial anisotropy below the noise level, and one must include in this noise the uncertain contribution from foreground contamination. If one has low detector noise and good spectral coverage one can determine the foreground contribution accurately, and thus there is little foreground uncertainty or noise. However if one cannot determine the foreground contribution accurately then the foreground noise is large and one should account for this in setting limits. This philosophy then leads us to reject both alternatives 1) and 2), which take no account of the added uncertainties due to foreground. Alternative 2) is especially dangerous since one not only underestimates the uncertainties but if the uncertainties in the foreground are large, one may also subtract away much of the MBR signal. This is illustrated for the SP91 experiment in fig. 3. The small spectral coverage of the experiment allows for very little discrimination between free-free emission and primordial anisotropy, and subtracting the best fit free-free signal therefore will lead to the subtraction of much of any primordial anisotropy present.

Of course, the problem with alternative 3) is determining which points are likely to be contaminated in a way which is unrelated to the amplitude of the primordial anisotropy at those points. If ones criterion for culling a data point has anything to do with the amplitude of the signal obtained at that data point, then one runs the risk of biasing the limits on primordial anisotropy. One proper way of culling data from multi-frequency experiments is to drop all of the data in patches where the signals which are linearly independent of primordial anisotropies are particularly high. Such linear combinations are found by solving equations analogous to eq. (12).

In many analyses of the SP91 experiment, only the 4th channel is used, essentially culling data from channels 1-3. One would hope that the rationale for this comes from the *a priori*

knowledge that experiments in this frequency range are most likely to be contaminated by free-free and synchrotron emission and therefore channel 4 will be the least contaminated. Unfortunately the signal in channels 1-3 is significantly greater than that in the 4th channel, which essentially meant that the data that had significant signal was dropped and that which had none was kept. Since the channels are so close in frequency the signals in channels 1-3 give us a strong indication that whatever is causing this signal, and it might be largely primordial anisotropy, should also be present in channel 4 since only an implausibly steep spectrum would give a negligible contribution to this highest frequency channel. We would therefore argue that the uncertainty in the foreground contamination of channel 4 is quite large and therefore any small upper limit on $Q_{\text{rms-ps}}$ derived from analyzing channel 4 while ignoring this foreground noise is unreliable.

One might argue that upper limits such as that of eq. (38) are too weak since the limit is weaker than that obtained by considering only channel 4 and ignoring foreground contamination. The rationale for this argument is that it is much more probable that the foreground will increase the observed signal rather than decrease it. We do not find these arguments very convincing. After all, given only 7 degrees of freedom the probability that χ^2 is less than half its expected value is 0.165, not a very small number. It is interesting to note that if, say, half the signal comes from primordial anisotropy and half the signal comes from foreground contamination, then a low χ^2 is much more likely to be produced by a cancellation of the foreground and the primordial anisotropy than by having the foreground and primordial contribution both be small. This is easily understood in terms of phase space arguments as follows. Consider two random variables X and Y both uniformly distributed in $[-1, 1]$ (see Figure 6). Thus the probability distribution in the X - Y plane is uniform in a square. The locus of possible outcomes which have a given value of $X + Y$ are just lines that cut through the square at a -45° , and if this sum is much smaller than 1 this line passes close to the origin, i.e. close to the diagonal $X = -Y$. Most of the length of such a diagonal line is not located near the origin where both X and Y are small but is rather closer to the corners where $X, Y \sim 1$ and there is strong cancellation. Thus if $X + Y$ is small it is more likely that $X \approx -Y$ rather than both X and Y being small. The same arguments works when a Gaussian distribution replaces a uniform distribution. To reiterate: cancellation is likely where the total signal is low even though it is unlikely in general. One may apply this to channel 4 of SP91 where we think there is evidence that the signal is low. If both foreground and primordial anisotropy are contributing to channels 1-3 we shouldn't be surprised to find significant cancellation between primordial and foreground contamination in channel 4.

In our analysis we have not made use of any other data outside of the microwave anisotropy data to limit the amount of foreground contamination. One should be careful with such procedures




FIG. 6. A plane in which X and Y are randomly distributed. Even if data constrains $|X + Y| < .05$ [the shaded region in the figure], neither X nor Y is necessarily small. In fact, in much of the shaded region, both $|X|$ and $|Y|$ are of order one.

as they often require bold extrapolations of the spectrum. In contrast, to marginalize we require only that the assumed form of the spectrum holds over a relatively narrow frequency region [i.e. there is a big difference between extrapolating a ν^{-3} spectrum from 408 MHz to 40 GHz and extrapolating the same ν^{-3} spectrum between 25 and 35 GHz]. While there is often reliable physics behind the extrapolations there are usually some caveats behind the applicability of these extrapolations. For example free-free or synchrotron self-absorption may severely decrease emission in the radio of sources which are bright in the microwave. Of course, if we have no idea about the spectra of the foreground we cannot marginalize the signal from these sources. Let us hope that sources with such problematic spectra do not exist in any abundance. In spite of the caveats we do take seriously these measurements. In many cases the data indicates that there should be no significant contamination, and in these cases marginalization may yield overly broad upper and lower limits.

In the future what must be hoped for is better data with low enough noise in enough different frequency channels to be able to fit all of the components to microwave brightness fluctuations. We would like to point out that the null space analysis of eq. (12) can be useful in choosing which channels to use for such experiments. Many of the very large upper limits in the Table are simply a result of the coefficients obtained by solving eq. (12). In other words, for some choices of channel frequencies the signal-to-noise is greatly reduced by projecting out the foregrounds while for other choices the reduction is less severe. By choosing the right frequencies one can optimize the signal-to-noise which is left after marginalization. Of course, instrumental considerations as well as considerations of atmospheric emission must also play a strong role in the choices of frequencies.

This work was supported in part by the DOE and NASA grant NAGW-2381 at Fermilab.

REFERENCES

- Bond, J.R., Efstathiou, G., Lubin, P.M., and Meinhold, P.R. 1991, PRL 66, 2179
- Brandt, W.N., Lawrence, C.R., Readhead, A.C.S., Pakianathan, J.N., and Fiola, T.M. 1993, CalTech-OVRO preprint # 12
- Gaier, T., Shuster, J., Gundersen, J., Koch, T., Meinhold, P., Sieffert, M., and Lubin, P., 1992, ApJ 398, L1
- Gundersen, J., *et al.* 1993, ApJ 413, L1
- Meinhold, P., *et al.* 1993, ApJ 409, L1
- Myers S.T. 1990, CalTech Ph.D. thesis
- Meinhold, P. and Lubin, P. 1991, ApJ 370, L11
- Schuster, J., Gaier, T., Gundersen, J., Meinhold, P., Koch, T., Sieffert, M., Wuensche, S.C., and Lubin, P. 1993, ApJ 412, L47
- Smoot, G.F., *et al.* 1992, ApJ 396, L1
- Srednicki, M., White, M., Scott, D., and Bunn, E.T. 1993, CfPA-93-TH-27 preprint